



## OPEN ACCESS

## EDITED BY

Farook Sattar,  
University of Victoria, Canada

## REVIEWED BY

Xiaoqiang Hua,  
National University of Defense Technology,  
China  
Mohamed Nabih,  
Bruno Kessler Foundation (FBK), Italy

## \*CORRESPONDENCE

Hiroshi Sato,  
✉ hrs.sato@ntt.com

RECEIVED 05 March 2025

ACCEPTED 11 June 2025

PUBLISHED 10 July 2025

## CITATION

Sato H, Ochiai T, Delcroix M, Moriya T,  
Ashihara T and Masumura R (2025) Generic  
speech enhancement with self-supervised  
representation space loss.  
*Front. Signal Process.* 5:1587969.  
doi: 10.3389/frsip.2025.1587969

## COPYRIGHT

© 2025 Sato, Ochiai, Delcroix, Moriya, Ashihara  
and Masumura. This is an open-access article  
distributed under the terms of the [Creative  
Commons Attribution License \(CC BY\)](#). The use,  
distribution or reproduction in other forums is  
permitted, provided the original author(s) and  
the copyright owner(s) are credited and that the  
original publication in this journal is cited, in  
accordance with accepted academic practice.  
No use, distribution or reproduction is  
permitted which does not comply with these  
terms.

# Generic speech enhancement with self-supervised representation space loss

Hiroshi Sato\*, Tsubasa Ochiai, Marc Delcroix, Takafumi Moriya,  
Takanori Ashihara and Ryo Masumura

NTT Corporation, Yokosuka-shi, Kanagawa, Japan

Single-channel speech enhancement is utilized in various tasks to mitigate the effect of interfering signals. Conventionally, to ensure the speech enhancement performs optimally, the speech enhancement has needed to be tuned for each task. Thus, generalizing speech enhancement models to unknown downstream tasks has been challenging. This study aims to construct a generic speech enhancement front-end that can improve the performance of back-ends to solve multiple downstream tasks. To this end, we propose a novel training criterion that minimizes the distance between the enhanced and the ground truth clean signal in the feature representation domain of self-supervised learning models. Since self-supervised learning feature representations effectively express high-level speech information useful for solving various downstream tasks, the proposal is expected to make speech enhancement models preserve such information. Experimental validation demonstrates that the proposal improves the performance of multiple speech tasks while maintaining the perceptual quality of the enhanced signal.

## KEYWORDS

self-supervised learning, loss function, SUPERB benchmark, signal denoising, speech enhancement, deep learning, speech recognition

## 1 Introduction

The recent advancements in machine learning technology have rapidly improved the level of machine comprehension of natural language. In our daily lives, natural language often serves as a medium for communication in the form of speech. To achieve machine understanding on the basis of such spontaneous speech and to enable natural responses to spoken language, studies have extensively explored various speech tasks such as automatic speech recognition (ASR), automatic speaker verification (ASV) (Yang et al., 2021; Tsai et al., 2022; Prabhavalkar et al., 2023; Huh et al., 2023; Wani et al., 2021). To apply these speech technologies in real-world environments, which often include various types of noise, speech enhancement (SE) is often introduced as the front-end that suppresses such noises as urban bustle, automobile sounds, household clatter from dishes, and the keystrokes and clicks of typing (Delcroix et al., 2018; Eskimez et al., 2018; O'Malley et al., 2021; Pandey et al., 2021; Chang X. et al., 2022; Masuyama et al., 2023; Lu et al., 2022; Fujita et al., 2024; Kandagatla and Potluri, 2020). It has been shown that an SE front-end can improve subsequent speech tasks such as ASR (Kinoshita et al., 2020), speech emotion recognition (Avila et al., 2018), and ASV (Shon et al., 2019). In this work, we call the speech enhancement module to enhance input signal for subsequent tasks as “front-end,” and the subsequent task as “back-end.”

Besides, the SE front-end is crucial for improving human listening. For instance, SE technology is widely applied in web conferences to remove background noise and facilitate smoother communication. Various research efforts have been devoted to improving the quality of speech recording (Reddy et al., 2020; Dubey et al., 2022; Dubey et al., 2024). Briefly stated, SE front-end is a general technique to improve automatic speech processing technologies and human auditory experiences. Hereafter, the back-end tasks related to machine-based speech processing, along with the task aimed at improving human listening, will be collectively referred to as *downstream tasks*.

We focus on single-channel SE that aims at extracting a clean speech signal from a monaural recording of noisy speech. Numerous methods have been proposed to date as single-channel SE techniques (Reddy et al., 2020; Dubey et al., 2022; 2024). SE is broadly divided into two categories: either in the time-frequency (TF) domain (Wang et al., 2019; Žmolíková et al., 2019; Hu et al., 2020; Xia et al., 2020; Hao et al., 2021; Zhao et al., 2022; Ju et al., 2023) or directly in the time domain (Luo and Mesgarani, 2019; Delcroix et al., 2020; Pandey and Wang, 2020; Défossez et al., 2020; Wang et al., 2021; Sato et al., 2024). Deep complex convolution recurrent network (DCCRN) is one of the major TF domain approaches (Hu et al., 2020). DCCRN enhances complex valued TF representation with the backbone network using a recurrent neural network (RNN) and convolutional neural network (CNN) and performs so well that it ranked first in the real-time-track in Interspeech 2020 Deep Noise Suppression (DNS) Challenge (Reddy et al., 2020). Conv-TasNet is a representative method of the time-domain approach, which directly enhances the input waveform to obtain the enhanced waveform (Luo and Mesgarani, 2019). Initially proposed for speech separation, Conv-TasNet performs superiorly to the ideal TF mask.

However, the improvements in SE performance measures such as signal-to-distortion ratio (SDR) do not necessarily mean that the SE method will improve the performance of downstream tasks when used as the front-end. It is reported that SE sometimes degrades back-end task performance despite the improvement in such performance measures as SDR (Yoshioka et al., 2015; Chen et al., 2018; Fujimoto and Kawai, 2019). This is because deep neural network (DNN)-based SE tends to generate “processing artifacts” due to non-linear transformations, which are detrimental to the subsequent downstream tasks (Iwamoto et al., 2022; Sato et al., 2022; 2021). Since such processing artifacts are unknown to downstream tasks during training, the negative effects of these distortions tend to outweigh the benefits of noise removal. In other words, there is a mismatch between the output of the SE whose training objective (e.g., SDR) attends to remove noise as much as possible while introducing artifacts and the input expected by subsequent tasks, which may be sensitive to unexpected artifacts.

Several methods have been investigated to reduce the mismatch between the SE front-end and the back-end. One approach is to train the back-end model with the enhanced signal, which is output by the SE front-end. This method allows the back-end model to learn, including the distortions present in the enhanced signal, thereby reducing the mismatch (Kinoshita et al., 2020). However, this method makes the back-end dependent on a specific front-end enhancement model, which compromises the system’s modularity, making it impossible to develop the front-end and

back-end independently. Additionally, this approach cannot be applied if the subsequent task involves using API services to solve back-end speech tasks, such as Google’s speech-to-text API (Speech to text, 2024), Open AI speech-to-text API (Speech-to-Text AI, 2024), or employing off-the-shelf black-box models since we cannot retrain them. Besides, this approach may not be applicable when the back-end model is extremely large, such as large audio-language models (LALMs), where retraining is not feasible.

Another approach is to integrate the SE front-end with the back-end task and jointly optimize them using the training criteria of the latter task (Gao et al., 2015; Wang and Wang, 2016; Menne et al., 2019; Subramanian et al., 2019). This can mitigate the mismatch between the two models. However, this method has a limitation that the SE front-end becomes specialized for a particular subsequent system, which may prevent it from fully performing when combined with different back-end systems. From a practical point of view, it is ideal to improve the performances of the SE front-end and back-end models separately. Requiring joint training every time one component is updated incurs significant operational costs, which is often not practical.

In this study, we aim to develop a single, generic SE model that is constructed independently of subsequent tasks and can be applied to various downstream applications, including automatic speech understanding technologies and human auditory experiences. The core idea is to train the SE model using a loss function that encompasses speech representations suitable for various downstream tasks. Recently, self-supervised learning (SSL) has been proposed and has proven effective in learning such generic representations. On the basis of this, we propose the SSL Mean Square Error (SSL-MSE) loss, which computes the loss within the SSL representation space. SSL, which is a research field that has gained significant attention in recent years, can learn useful representations without the need for artificially labeled data (Mohamed et al., 2022; Liu et al., 2021; Jaiswal et al., 2020). The proposed SSL-MSE loss is a training criterion to minimize the distance between the ground truth clean signal and enhanced signal in the feature domain that is extracted by the pre-trained SSL model. Since SSL models have been shown to learn effective representations for many downstream tasks, they can capture not only acoustic but also higher-level information such as phonetic or semantic information (Pasad et al., 2021; Dunbar et al., 2021; Hsieh et al., 2021). Thus, we expect that creating a loss term in the SSL domain could guide SE training to preserve or enhance various levels of characteristics of the speech signals that may be required for high-level downstream tasks.

This approach enables not only a single SE model to be applied to multiple tasks but also the SE front-end and the back-end to be independently improved in solving each task, thereby maintaining modularity and allowing each technology to be independently maintained and developed. Additionally, an SE model will be able to be constructed that can improve the performance of downstream tasks that are black-box and cannot be retrained, as well as tasks with a very large number of parameters and for which retraining is impractical.

In this work, we evaluate the effectiveness of the proposal with various back-end tasks. We evaluated speech recognition, speaker verification, and intent classification performance using the SSL model in the pipeline, as well as a speech recognition model using an off-the-shelf Whisper model and a human listening task, i.e., measuring the objective measures of perceptual quality of the

enhancement signal by perceptual evaluation of speech quality (PESQ) (Rix et al., 2001) and deep noise suppression mean opinion score (DNSMOS) (Reddy et al., 2022).

The main contributions of this paper can be summarized as follows:

1. We propose a novel training loss that minimizes the distance between enhanced and clean speech signals in the SSL representation space. The proposed SSL-MSE loss transfers the generalizability of SSL models over various downstream tasks into SE models.
2. We perform an extensive evaluation on multiple downstream tasks, showing that models trained with the proposed SSL-MSE loss outperform vanilla SE models as well as models trained on different representation spaces such as log mel-filter bank (LMFB) or ASR loss.

While our approach shares the idea of utilizing SSL representations with the prior work proposed by Hsieh et al. (Hsieh et al., 2021), there are key differences in both the formulation and objectives. Hsieh et al. primarily aim to improve perceptual quality for human listeners and proposed phonefortified perceptual loss (PFPL) defined on the final-layer output of wav2vec 2.0. On the other hand, this study focused on building a generic SE model that is effective across a wide range of downstream tasks. To obtain such versatility, we proposed the SSL-MSE loss that incorporates weighted representations from the latter half of the SSL layers. Additionally, to show the versatility for various downstream tasks, we have conducted a comprehensive experiments using a variety of SSL models and downstream task evaluations. These distinctions in design and effect underline the novelty and utility of our approach. This paper extends our previous short conference paper, where we introduced the basic concept of SSL-MSE and showed that SSL-MSE loss could help downstream models that utilize SSL models for feature extraction (Sato et al., 2023). In this work, we provide a more detailed explanation of SSL-MSE, adding more context. In addition, we include new experimental results showing that it can also benefit models that do not rely on SSL features, thereby achieving a truly generic model. To evaluate the generalizability over downstream tasks, we extend experimental validations to include objective measures of perceptual quality, such as PESQ (Rix et al., 2001) and DNSMOS (Reddy et al., 2022), which are known to correlate with human subjective evaluation, as well as ASR performance evaluation using the Whisper model (Radford et al., 2023). Furthermore, we examine a combination of SSL-MSE training with an observation adding (OA) post-processing technique to further reduce the effect of the distortion generated by the SE.

The rest of the paper is organized as follows. In Section 2, we will discuss the related works. In Section 3, we explain the conventional method, and in Section 4, we explain the proposed approach. In Section 5, we detail the experimental validation. Section 6 concludes the paper.

## 2 Related works

### 2.1 SSL models

SSL in speech is a technology for learning powerful representations. For speech SSL models, various techniques have

been proposed (Liu et al., 2020a; Liu et al., 2020b; Schneider et al., 2019; Baevski et al., 2020a; Chang H.-J. et al., 2022; Baevski et al., 2020b; Hsu et al., 2021; Chen et al., 2022; 2023; Yadav et al., 2024; Shi et al., 2024), among which wav2vec 2.0 (Baevski et al., 2020b), HuBERT (Hsu et al., 2021), and WavLM (Chen et al., 2022) are some of the most widely applied approaches.

The SSL model is trained to extract a time series of feature representation from a single-channel audio, which is useful for many types of downstream tasks that can be learned on the representation with a limited amount of paired data. The task of training the upstream SSL model is called the pre-text task, which differs for each type of SSL model.

For example, wav2vec 2.0 (Baevski et al., 2020b) masks parts of the input speech in the latent space and solves a contrastive prediction task over quantized representations, encouraging the model to learn the structure and patterns of speech without labeled data. Another key aspect of wav2vec 2.0 is its use of quantization. After initial processing, the audio is transformed into quantized representations by mapping the continuous audio features into a finite set of codebook entries. Quantization enables the model to represent audio with a smaller, more manageable set of discrete units, which captures essential information about the speech content while reducing redundancy. This discrete representation is used as the target for the model's predictions, allowing it to focus on the fundamental components of speech rather than unnecessary detail. HuBERT (Hsu et al., 2021) and WavLM (Chen et al., 2022) are trained using the BERT-like masked prediction task (Devlin et al., 2019) on the target label generated in an offline clustering step. By masking of the input, HuBERT is trained to predict these hidden units for the masked frames, helping it capture essential patterns in speech structure and phonetic information. Of particular note, WavLM is made robust to noise and interference speakers through introducing a denoising task into the pre-text task. Specifically, DNS noise and interfering speech are added to the input audio, and WavLM is trained to ignore these signals. Although a prior study showed that WavLM was relatively robust to noise (Masuyama et al., 2023), Chung et al. revealed that the noise robustness of SSL models could be further improved by introducing a single-channel SE front-end (Chang X. et al., 2022).

Another approaches have been proposed for improving the noise robustness of SSL system by performing speech enhancement directly in the SSL feature domain (Ali et al., 2023; 2022). As a result of the training on the pre-text task, SSL upstream models function as a universal feature extractor that can generally be applied for various downstream tasks (Yang et al., 2021; Tsai et al., 2022). Note that the feature extracted from the SSL model is calculated as the weighted sum over the outputs from the transformer layers of SSL models where different weights are used for different downstream tasks. This reveals that SSL models can capture rich speech representation through their layers. The proposed SSL-MSE aims to transfer the generalizability of the SSL model into an SE model by using the loss term calculated on the SSL representations.

### 2.2 Processing artifacts

Another deeply related research topic is the processing artifacts generated by SE. DNN-based SE generates processing artifacts that

are detrimental to subsequent tasks. As mentioned in Introduction, joint training is one way to mitigate artifacts, but it implies creating SE specific for each downstream task, which is not practical. A promising approach to mitigate the mismatch caused by the processing artifact is OA post-processing, which interpolates the enhanced and observed signals as the input signal to the back-end (Ochiai et al., 2024). OA improves the signal-to-artifact ratio instead of increasing the signal-to-noise ratio (SNR). The appropriate ratio of OA improves back-end ASR performance since artifacts tend to be more detrimental than noise to ASR. In this work, we investigate the combination of the proposed SSL-MSE and the OA approach.

## 3 Fundamental system overview

### 3.1 Self-supervised learning

There are two major ways to apply SSL models to downstream tasks, either as a fixed feature extractor, or permitting retraining of their parameters (Mohamed et al., 2022). Following previous studies (Yang et al., 2021; Chang X. et al., 2022; Masuyama et al., 2023), we adopted the former approach, i.e., freezing the parameters of the SSL model  $\theta_{ssl}$  during the training of the downstream models and using the SSL upstream model as a feature extractor, because SSL models are usually very large and thus fine-tuning the SSL models for every downstream task is computationally too intensive. The downstream task-specific additional layers are relatively small DNNs, which we call the downstream model.

Formally, we can write the feature extraction process with the SSL model as shown in Equation 1 (Yang et al., 2021):

$$F_{1:N} = \text{SSL}(\mathbf{x}; \theta_{ssl}), \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^T$  denotes the monaural input signal in raw waveform with  $T$  samples,  $F_n \in \mathbb{R}^{D \times T'}$  denotes the time series of extracted features obtained from the  $n$ -th layer of the SSL model, and  $F_{i:j}$  denotes  $F_{i:j} = [F_i, F_{i+1}, \dots, F_j]$ .  $N$  denotes the number of layers of the SSL model, and  $\theta_{ssl}$  denotes the learnable parameters of the SSL model.  $T'$  and  $D$  denote the number of frames and dimensions of the extracted features, respectively. The learnable parameters,  $\theta_{ssl}$  are trained with the pre-text task.

An effective way to apply the SSL features across various types of downstream tasks is to use the weighted sum of the embeddings from different layers in the SSL model as the input feature of the downstream model (Chen et al., 2022; Yang et al., 2021). The process of the downstream model for a task  $\tau$  can be written as shown in Equations 2, 3:

$$F^\tau = \text{WS}(F_{1:N}; \theta_w^\tau) \equiv \sum_{n=1}^N w_n^\tau F_n, \quad (2)$$

$$\hat{I}^\tau = \text{DS}^\tau(F^\tau; \theta_{ds}^\tau), \quad (3)$$

where  $\text{WS}(\cdot)$  is the weighted sum function,  $\theta_w^\tau = [w_1^\tau, \dots, w_N^\tau]$  are the weights for latent representations obtained from each layer, which are learnable parameters, and  $\hat{I}^\tau$  is the estimation result attained by the downstream model.  $\text{DS}^\tau(\cdot)$  denotes the downstream model for the task  $\tau$  whose learnable parameters are  $\theta_{ds}^\tau$ . Since SSL models capture varying levels of representation across layers, the latent representation from each layer  $F_n$  retains

information at different granularities. By using the weighted sum of these representations as input to the downstream model, the downstream model is expected to be able to effectively solve a variety of tasks. The learnable parameters  $(\theta_w^\tau, \theta_{ds}^\tau)$  are jointly optimized by using task-specific paired data  $(I^\tau, \hat{I}^\tau)$  where  $I^\tau$  is the ground truth label, while the upstream model  $\theta_{ssl}$  is frozen.

### 3.2 Speech enhancement

We investigate single-channel neural-based SE as the front-end. Let  $\mathbf{y} \in \mathbb{R}^T$  and  $\hat{\mathbf{x}} \in \mathbb{R}^T$  be the noisy observations and enhanced signals, respectively. The enhancement process can be denoted as shown in Equation 4:

$$\hat{\mathbf{x}} = \text{SE}(\mathbf{y}; \theta_{se}). \quad (4)$$

$\theta_{se}$  denotes the learnable parameters of the SE model. To train the SE model, paired data of noisy observation  $\mathbf{y}$  and ground truth clean source  $\mathbf{x} \in \mathbb{R}^T$  are commonly prepared. The SE model is commonly optimized by minimizing the distance between enhanced signal  $\hat{\mathbf{x}}$  and the ground truth clean source  $\mathbf{x}$ . Specifically, we adopt the scale-dependent SNR loss  $\mathcal{L}_{\text{SNR}}(\cdot)$  as the distance measure, which is defined as shown in Equation 5:

$$\mathcal{L}_{\text{SNR}} = -10 \log_{10} \frac{\|\hat{\mathbf{x}}\|^2}{\|\mathbf{x} - \hat{\mathbf{x}}\|^2}. \quad (5)$$

Since the clean source is not available for real-recorded noisy speech, simulated noisy speech  $\mathbf{y}$  is generated on the basis of the speech  $\mathbf{x}$  and noise recordings for SE training (Vincent et al., 2006).

### 3.3 Combination of SE front-end with back-end tasks

The introduction of an SE front-end can improve the performance of subsequent tasks. The subsequent task can either be 1) a speech processing task with SSL upstream and downstream models, 2) a speech processing task without SSL models, or 3) a human listening task. The overview of the combination of SE front-end with back-end tasks is shown in Figure 1. Let the feature extraction process be  $\text{FE}(\cdot)$ , the general pipeline to solve a downstream task with SE can be expressed as shown in Equation 6:

$$\hat{I}^\tau = \text{DS}^\tau(\text{FE}(\text{SE}(\mathbf{y}; \theta_{se}); \theta_{fe}); \theta_{ds}^\tau). \quad (6)$$

where  $\theta_{fe}$  denotes the learnable parameters of the feature extractor.

#### 3.3.1 Back-end with SSL models

For a speech understanding task with SSL models, the feature extractor in Equation 6 is an SSL upstream model and weighted sum function. The whole pipeline can be expressed as shown in Equation 7 (Yang et al., 2021):

$$\hat{I}^\tau = \text{DS}^\tau(\text{WS}(\text{SSL}(\text{SE}(\mathbf{y}; \theta_{se}); \theta_{ssl}); \theta_w^\tau); \theta_{ds}^\tau). \quad (7)$$

Previous research reports that to benefit from SE, it is necessary to fine-tune SE that is trained independently from the SSL pipeline (Masuyama et al., 2023) to mitigate the mismatch between the front-end and the back-end. Specifically, to mitigate this mismatch,






FIGURE 1

Overview of the combination of SE front-end with back-end tasks. The backend task can either be 1) a speech processing task with SSL upstream and downstream models, 2) a speech processing task without SSL models, or 3) a human listening task.

$(\theta_{se}, \theta_w, \theta_{ds})$  are jointly optimized by using downstream task paired data while freezing SSL model  $\theta_{ssl}$  (Masuyama et al., 2023). However, the fine-tuning of the downstream task makes the SE front-end task-specific, making it infeasible to share a single general SE model for multiple downstream tasks, thus incurring development costs. Besides, when considering more than a few downstream tasks, optimizing the SE front-end and back-end model jointly for every downstream task may be computationally demanding.

### 3.3.2 Back-end without SSL models

For the back-end tasks without SSL models, the feature extractor in Equation 6 is typically an LMFB feature extractor. The whole processing pipeline can be expressed as shown in Equation 8:

$$\hat{\mathbf{t}} = \text{DS}^T(\text{LMFB}(\text{SE}(\mathbf{y}; \theta_{se})); \theta_{ds}). \quad (8)$$

For systems that do not use SSL, the mismatch between the SE front-end and back-end is also a challenge. As a way to mitigate the mismatch, LMFB feature domain loss has been proposed that calculates the loss term on the LMFB feature domain (Wang et al., 2020). Preceding works show the effectiveness of LMFB feature domain loss in improving the back-end ASR task with SE. In this work, we examine the multitask loss of LMFB feature domain loss and SNR loss. Formally, the multitask loss  $\mathcal{L}_{\text{LMFB-MT}}$  is expressed as shown in Equations 9, 10 (Wang et al., 2020):

$$\mathcal{L}_{\text{LMFB-MT}} = \mathcal{L}_{\text{LMFB}}(\hat{\mathbf{x}}, \mathbf{x}) + \alpha \mathcal{L}_{\text{SNR}}(\hat{\mathbf{x}}, \mathbf{x}), \quad (9)$$

$$\mathcal{L}_{\text{LMFB}} = \|\text{LMFB}(\hat{\mathbf{x}}) - \text{LMFB}(\mathbf{x})\|_F^2 / M \quad (10)$$

where  $\|\cdot\|_F$  represents Frobenius norm and  $\text{LMFB}(\cdot)$  represents the LMFB feature extraction.  $\alpha$  represents a multitask weight.  $M$  represents the number of elements of LMFB feature  $\text{LMFB}(\mathbf{x})$ .

Another way to mitigate the mismatch is to train SE on the loss function that minimizes the distance between enhanced and clean signals on the output of the back-end task. Considering ASR as an example and letting  $\text{ASR}(\cdot)$  be the process of the ASR, the ASR result is obtained from a noisy observation via SE as shown in Equation 11:

$$\hat{\mathbf{w}} = \text{ASR}(\text{LMFB}(\text{SE}(\mathbf{y}; \theta_{se})); \theta_{asr}), \quad (11)$$

where  $\theta_{asr}$  and  $\hat{\mathbf{w}}$  indicate learnable parameters of the ASR model and ASR result, respectively. The ASR output basis loss minimizes the distance between the ASR result of the enhanced signal  $\hat{\mathbf{w}}$  and ground truth transcription  $\mathbf{w}$  to optimize the SE model  $\theta_{se}$  while freezing the ASR model  $\theta_{asr}$ . We hereafter refer to this training

criterion as ASR loss. We examine the SE training with the multitask loss of the ASR loss and SNR loss as shown in Equations 12, 13:

$$\mathcal{L}_{\text{ASR-MT}} = \mathcal{L}_{\text{ASR}}(\hat{\mathbf{w}}, \mathbf{w}) + \alpha \mathcal{L}_{\text{SNR}}(\hat{\mathbf{x}}, \mathbf{x}), \quad (12)$$

$$\mathcal{L}_{\text{ASR}} = \lambda \mathcal{L}_{\text{ctc}}(\hat{\mathbf{w}}, \mathbf{w}) + (1 - \lambda) \mathcal{L}_{\text{att}}(\hat{\mathbf{w}}, \mathbf{w}), \quad (13)$$

where  $\alpha$  is the multitask weight controlling the impact of SNR loss.  $\mathcal{L}_{\text{ctc}}$  and  $\mathcal{L}_{\text{att}}$  optimize connectionist temporal classification (CTC) (Graves et al., 2006) decoder and attention-based decoder, respectively. Note that  $\mathcal{L}_{\text{att}}$  is a cross-entropy loss.  $\lambda$  is a hyperparameter controlling the balance of CTC and attention loss.

### 3.3.3 Human perception as back-end task

We consider human perception as a back-end task. In this case, there is no feature extraction  $\text{FE}(\cdot)$  or downstream model  $\text{DS}^T(\cdot)$ , as the goal is to produce an enhanced signal, and thus the whole system can be expressed as Equation 4.

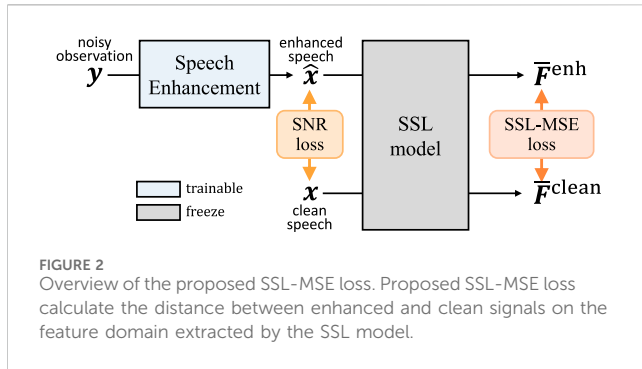
## 4 Proposed methods

To realize an SE that can be applicable to multiple tasks, we propose SSL-MSE. Figure 2 shows the general framework of the proposed SSL-MSE training. SSL-MSE loss calculates the distance between enhanced and clean signals on the SSL model's feature domain. Since SSL models have been shown to learn effective representations for many downstream tasks, SSL-MSE can make SE more optimal for the front-end of multiple tasks. More specifically, SSL models are reported to capture not only acoustic but also higher-level information such as phonetic or semantic information (Pasad et al., 2021; Dunbar et al., 2021; Hsieh et al., 2021), and thus a loss term in the SSL domain could guide SE training to preserve or enhance various levels of characteristics of the speech signals that may be required for high-level downstream tasks. Formally, SSL-MSE  $\mathcal{L}_{\text{SSL}}$  between enhanced speech  $\hat{\mathbf{x}} = \text{SE}(\mathbf{y}; \theta_{se})$  and clean source  $\mathbf{x}$  is calculated as the MSE between SSL features extracted from these signals as shown in Equations 14–18:

$$\mathcal{L}_{\text{SSL}} = \|\bar{\mathbf{F}}^{\text{enh}} - \bar{\mathbf{F}}^{\text{clean}}\|_F^2 / (DT') \quad (14)$$

$$\bar{\mathbf{F}}^{\text{enh}} = \sum_{n=1}^N \tilde{w}_n \mathbf{F}_n^{\text{enh}}, \quad (15)$$

$$\bar{\mathbf{F}}^{\text{clean}} = \sum_{n=1}^N \tilde{w}_n \mathbf{F}_n^{\text{clean}}, \quad (16)$$



$$\mathbf{F}_{1:N}^{\text{enh}} = \text{SSL}(\text{SE}(\mathbf{y}; \theta_{\text{se}}); \theta_{\text{ssl}}), \quad (17)$$

$$\mathbf{F}_{1:N}^{\text{clean}} = \text{SSL}(\mathbf{x}; \theta_{\text{ssl}}), \quad (18)$$

where  $\tilde{\mathbf{w}} = [\tilde{w}_1, \dots, \tilde{w}_N]$  represents the layer weight for the latent feature obtained from each layer of the SSL model. The layer weight is a hyperparameter that controls which layer output SSL-MSE focuses on. Since SSL models capture different levels of representations in their different layers, the latent representation from each layer retains information at different levels of granularity. Thus, applying SSL loss over multiple layers may be advantageous as it ensures that the SE module captures necessary information. Specifically, we adopt a weight that uniformly emphasizes the features outputted from the latter half of the layers to preserve higher-level representations, as the lower-level ones are expected to be captured in combination with SNR loss. Formally, the weight is expressed as shown in Equations 19, 20:

$$\tilde{w}_1, \dots, \tilde{w}_{\lfloor 2N/2 \rfloor} = 0, \quad (19)$$

$$\tilde{w}_{\lfloor 2N/2 \rfloor + 1}, \dots, \tilde{w}_N = \frac{1}{N - \lfloor 2N/2 \rfloor}, \quad (20)$$

where  $\lfloor \cdot \rfloor$  represents the floor function. Although it is possible to calculate the loss on each layer and sum the losses instead of calculating the loss on the weighted sum of the representations, we adopted the latter way to maintain consistency with the feature extraction process of the SSL pipeline as in Equation 7.

The model parameter  $\theta_{\text{se}}$  is optimized by minimizing the multitask loss function  $\mathcal{L}_{\text{SSL-MT}}$  expressed as shown in Equation 21:

$$\mathcal{L}_{\text{SSL-MT}} = \mathcal{L}_{\text{SSL}} + \alpha \mathcal{L}_{\text{SNR}}, \quad (21)$$

where  $\alpha$  denotes the multitask weight for the SNR loss. Note that the parameters of the SSL model  $\theta_{\text{ssl}}$  are frozen during the SSL-MSE training.

By developing an SE model that is trained independently from the back-end tasks and is applicable to multiple tasks, we can independently improve both the SE front-end and the back-end systems for each specific task. This approach preserves modularity, allowing each component to be independently maintained and developed. Furthermore, this methodology not only facilitates the use of a single SE model across various tasks but also enables SE models to be constructed that can improve the performance of downstream tasks, even when those tasks are black-box systems that cannot be retrained. It is also beneficial for tasks with a very large number of parameters, where retraining would be impractical.

To further reduce the mismatch between SE front-end and the back-end, we investigate the combination of SSL-MSE training with OA post-processing (Ochiai et al., 2024). OA post-processing interpolates enhanced and input signals in the waveform signal domain as the final output of the SE. The distortion generated by the SE can be decreased by adding the observation signal, which can also mitigate the mismatch between front-end and back-end. With OA, the input of the back-end task  $\tilde{\mathbf{x}}$  is expressed as shown in Equation 22:

$$\tilde{\mathbf{x}} = \beta \mathbf{y} + (1 - \beta) \hat{\mathbf{x}}, \quad (22)$$

where  $\beta$  is the OA ratio. Note that OA is implemented as a post-processing step for the enhanced signal, and thus depending on the downstream task, it is possible to modify the adding ratio of OA or omit its processing.

## 5 Experiments

### 5.1 Experimental setup

#### 5.1.1 Speech enhancement front-end

The SE model was trained on simulated mixtures of speech and noise data. We used LibriSpeech (Panayotov et al., 2015) for speech recordings and DNS noise (Reddy et al., 2021) for noise recordings sampled at 16 kHz. The number of noisy observations were 100,000 and 5,000 for the training and development sets, respectively. The noise was added at SNR values randomly sampled from  $-3$ – $20$  dB. We chose DNS noise to create a fair comparison of the SSL pipeline with/without the SE module, as DNS noise is also used in training the WavLM models. We adopted Conv-TasNet as the SE front-end module, which converts noisy raw-waveform audio into enhanced raw-waveforms in a time-domain, end-to-end processing manner (Luo and Mesgarani, 2019). In accordance with the setup adopted in (Thakker et al., 2022), we set the hyperparameters to  $N = 4096$ ,  $L = 320$ ,  $B = 256$ ,  $R = 4$ ,  $X = 8$ ,  $H = 512$  and  $P = 3$  following the notation in (Luo and Mesgarani, 2019). The SSL-MSE loss was introduced after pre-training the model with conventional SNR loss to speed up the conversion. The initial learning rate for the pre-training with SNR loss was set to  $5e-4$  and was multiplied by  $3/4$  if the loss on the development set did not decrease for 2 epochs. For optimization, we adopted the Adam optimizer (Kingma and Ba, 2015). The models were pre-trained for 100 epochs. If we apply OA, we adopt adding ratio  $\beta$  of 0.1 and 0.5.

#### 5.1.2 Multitask loss

We investigate three types of multitask losses: SSL-MSE loss, ASR loss, and LMF loss. Each multitask loss training is applied as the fine-tuning process: we first pretrained the SE model with SNR loss and then fine-tuned it with SSL-MSE loss with the initial learning rate of  $1e-4$  and for up to 50 epochs. We tested the SNR weight  $\alpha$  within  $\{0, 0.0001, 0.001, 0.01, 0.1, 1, 10\}$ . The comparison is shown in Figure 3.

SSL-MSE loss: To calculate SSL-MSE loss in Equation 21, we tested four off-the-shelf SSL models: WavLM BASE+, WavLM BARGE, wav2vec 2.0 BASE, and Hubert BASE models. SSL models were frozen

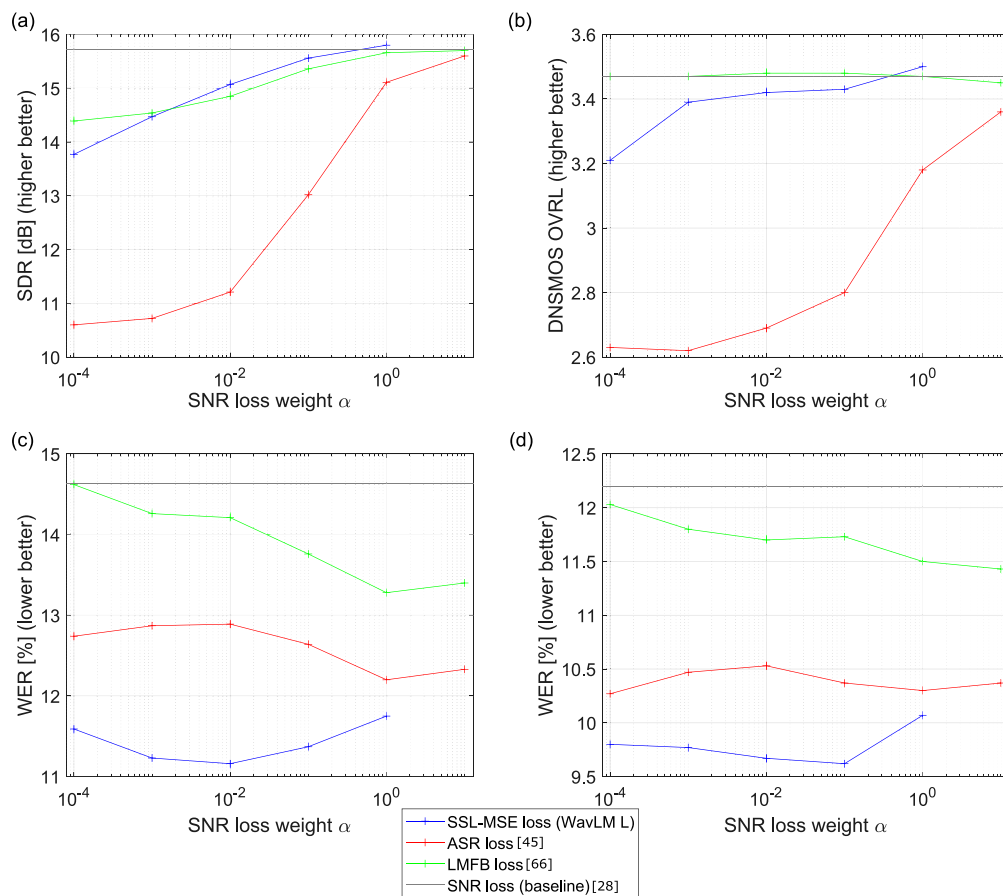


FIGURE 3

Comparison of SSL-MSE loss with ASR loss and LMFB loss joint training. The X-axis represents the multitask loss weight  $\alpha$  in the joint training with each loss, and the Y-axis represents the performance in each evaluation metric. For SDR [dB] DNSMOS OVRL, higher is better, and for WER [%], lower is better. (a) Speech enhancement performance. (b) DNSMOS result. (c) SSL downstream ASR with SE. (d) Whisper ASR with SE.

during SSL-MSE training. WavLM BASE+ and WavLM BARGE share the same training data; however, WavLM Large has more than three times the number of parameters compared to WavLM Base+. Wav2Vec 2.0 and HuBERT are not trained to be robust to noise; their pre-training procedures do not involve the explicit addition of noise. In contrast, WavLM BASE+ and WavLM LARGE (Chen et al., 2022) are pre-trained with noise augmentation to enhance robustness.

**ASR loss:** ASR loss in Equation 12 is calculated by a pre-trained transformer-based hybrid CTC/Attention ASR model (Watanabe et al., 2017; Karita et al., 2019), which consists of 12 conformer encoder blocks and 6 transformer decoder blocks. The model was trained following the LibriSpeech training recipe from ESPnet (Watanabe et al., 2018), an open-source toolkit for end-to-end speech processing.  $\lambda$  in Equation 12 was set as 0.3. The ASR model is frozen during the ASR loss training.

**LMFB loss:** For calculating LMFB loss in Equation 9 we adopt LMFB feature extractor of 80 mel bins, 400 sample window length, and 200 sample window shift.

### 5.1.3 Evaluation details

To evaluate the generalizability of the SE model trained with SSL-MSE loss, we evaluated the performance of SE in terms of the

following aspects: 1) SDR (Le Roux et al., 2019), 2) objective measures of subjective quality, 3) the performance of SSL downstream tasks, and 4) Whisper ASR performance. Objective measures of subjective quality is evaluated on PESQ and DNSMOS P.835 (Reddy et al., 2022). SDR, PESQ, and DNSMOS are evaluated on the 3,000 simulated mixtures of LibriSpeech speech recordings and DNS noise where SNR values are randomly selected from 0 to 10 dB.

**SDR:** SDR is evaluated on the 3,000 simulated mixtures of LibriSpeech speech recordings and DNS noise where SNR values are randomly selected from 0 to 10 dB.

**Objective measures of perceptual quality:** Perceptual quality is evaluated using objective measures, PESQ and DNSMOS P.835 (Reddy et al., 2022), which are known to correlate with human subjective evaluations using the same dataset as SDR evaluation.

**SSL downstream tasks:** To evaluate the performance in SSL downstream tasks, we prepared the SSL pipeline to solve three downstream tasks (automatic speech recognition (ASR), automatic speaker verification (ASV), and intent classification (IC)) following the procedure explained in Section 2.1. These performance are evaluated in terms of word error rate (WER), equal error rate (EER), and accuracy (Acc), respective. We implemented the SSL pipeline using the S3PRL toolkit (Yang

TABLE 1 Performance evaluation of the conventional and proposed systems.

	SE method or its training method	SSL model for SSL-MSE loss	OA ratio	SDR↑ [dB]	Human perception			SSL downstream tasks			Whisper	
					PESQ↑	DNSMOS			ASR	ASV	IC	ASR
						SIG↑	BAK↑	OVRL↑				
(a1)	clean			-	-	4.02	4.02	3.57	5.6	4.4	98.8	3.9
(a2)	no process			3.7	1.27	2.88	1.86	1.99	17.1	10.9	67.4	9.3
(b1)	SNR loss (Luo and Mesgarani, 2019)	N/A	N/A	<b>15.7</b>	2.27	3.73	<b>4.21</b>	<b>3.47</b>	14.6	8.6	84.8	12.2
(b2)		N/A	0.1	14.5	2.01	3.82	2.88	2.93	12.2	7.8	87.8	10.3
(c1)	SSL-MSE loss	WavLM B+	N/A	<b>15.8</b>	<b>2.35</b>	3.74	4.18	<b>3.46</b>	12.9	8.1	87.7	11.1
(c2)		WavLM L	N/A	15.6	<b>2.34</b>	3.82	3.94	3.42	11.3	7.4	88.6	9.6
(c3)		WavLM L	0.1	14.3	1.97	<b>3.85</b>	2.83	2.92	<b>10.9</b>	<b>7.2</b>	<b>88.7</b>	<b>9.1</b>

For the evaluation of SSL downstream tasks, we used the SSL pipeline with WavLM BASE+ model. The column SSL-MSE loss indicates the type of SSL model used for calculating SSL-MSE loss. WavLM B+ stands for WavLM BASE+, and WavLM L stands for WavLM LARGE. The SNR multitask loss weight  $\alpha$  for SSL-MSE loss training is fixed to 0.1. Values in bold indicate the results of t-tests conducted at a 5% significance level, showing no significant difference from the highest-performing results. N/A indicates not applicable for each term.

et al., 2021; Github, 2024). As the SSL model for the inference, we tested WavLM BASE+, WavLM LARGE, wav2vec 2.0 BASE, and Hubert BASE. The performance of SSL pipeline was evaluated on the noisy version of the SUPERB test sets for each task by adding DNS noise to the original recordings at SNR values randomly sampled from 0 to 10 dB. Since DNS noise contains a wide variety of the noise types and we use different samples for training and testing, we can verify robustness to unseen noise although we adopted DNS noise in both training and evaluation.

We tested two setups for downstream model training: 1) *official* setup where downstream models were trained with SUPERB official training data (Yang et al., 2021) that is relatively “clean” speech data, in accordance with the S3PRL SUPERB recipe (s3p, Cited October 7 2024), and 2) *noise-robust* setup where DNS noise was added to SUPERB official downstream training data at SNR values randomly sampled from −3–20 dB. The results gained in the noise-robust setup are shown in Table 3; for the other experiments, we adopt the official setup. The SE front-end was not applied while training the downstream model in noise-robust setups and in official setups. The downstream models were prepared for each downstream task for each SSL model.

**Whisper ASR:** Whisper ASR performance is measured with an off-the-shelf Whisper MEDIUM model on the simulated mixture of LibriSpeech speech recordings and DNS noise at the SNR values randomly sampled from 0 to 10 dB. The Whisper ASR model is only used for evaluation and not for ASR-loss training. Note that the Whisper model uses LMFB as a feature extraction and does not use SSL models.

## 5.2 Results and discussion

### 5.2.1 The effect of SSL-MSE

Table 1 compares the performances of SE models trained with conventional SNR loss training and proposed SSL-MSE loss. The

SNR multitask loss weight  $\alpha$  is fixed to 0.1. Comparing (a1) and (a2), the addition of noise to the test sets has a substantial impact on downstream performance.

For instance, the Whisper ASR system exhibits a WER degradation from 3.9% in (a1) to 9.3% in (a2) due to noise addition. This finding aligns with the observations in the original Whisper paper (Radford et al., 2023), particularly in Figure 5, which demonstrates that although Whisper models are trained to be robust to noise, they are not entirely impervious to its influence. The introduction of the baseline SE front-end (b) improves each performance metric, except Whisper ASR WER. Since the Whisper ASR model itself is trained to be robust to the noise, the negative effect of the mismatch between the SE front-end and back-end seems to outweigh the positive effect of the reduction of the noise. The systems (c1) and (c2) shows the effect of introducing the SSL-MSE loss, with a different SSL model used for calculating SSL-MSE loss, i.e., WavLM BASE PLUS (c1) and WavLM LARGE (c2). As for SDR and objective measures of subjective quality, the proposed systems (c1) and (c2) perform equivalently to the baseline system (b). More specifically, the proposed systems improve PESQ and DNSMOS SIG values and degrade DNSMOS BAK and DNSMOS OVRL values. The increase in the BAK value at the expense of the SIG value suggests that while the noise suppression effect is somewhat diminished, the harmful distortion that affects subsequent processing is reduced. The performances of SSL downstream tasks and Whisper ASR were substantially improved by introducing SSL-MSE loss. The improvement is larger when SSL-MSE loss is calculated with the WavLM LARGE model shown in (c2). The effect of the type of SSL model on the performance will be discussed in detail in Section 5.2.3. Compared with the baseline system (b), the proposed system (c2) improved the WER of the SSL downstream ASR task by 22%, the EER of the ASV task by 14%, and the Acc of the IC task by 3.8%. Based on the observation above, the introduction of SSL-MSE loss improves the performance as the front-end of the subsequent speech



**TABLE 2** Performance variations due to the combination of SSL model for inference and SSL model for training with SSL-MSE loss. w2v2 B, Hubert B, WavLM B+, and WavLM L stands for wav2vec 2.0 Base, Hubert Base, WavLM Base+, and WavLM Large, respectively. The SNR multitask loss weight  $\alpha$  for SSL-MSE loss training is fixed to 0.1.

	SSL model for inference	SE	SSL model for SSL-MSE loss	ASR	ASV	IC
				WER↓ [%]	EER↓ [%]	Acc↑ [%]
(a2)	WavLM B+			17.1	10.9	67.4
(b)	WavLM B+	✓	N/A	14.6	8.6	84.8
(d1)	WavLM B+	✓	w2v2 B	14.4	8.6	85.4
(d2)	WavLM B+	✓	Hubert B	14.7	8.7	86.1
(c1)	WavLM B+	✓	WavLM B+	12.9	8.1	87.7
(c2)	WavLM B+	✓	WavLM L	<b>11.3</b>	<b>7.4</b>	<b>88.6</b>
(e1)	w2v2 B			36.8	19.0	41.0
(e2)	w2v2 B	✓	N/A	18.8	12.2	70.9
(e3)	w2v2 B	✓	WavLM B+	17.5	11.3	71.3
(e4)	w2v2 B	✓	WavLM L	<b>14.9</b>	<b>11.0</b>	<b>75.9</b>
(e5)	w2v2 B	✓	w2v2 B	18.5	12.5	71.2
(f1)	Hubert B			31.7	15.5	57.8
(f2)	Hubert B	✓	N/A	18.5	10.6	81.9
(f3)	Hubert B	✓	WavLM B+	16.9	9.7	82.6
(f4)	Hubert B	✓	WavLM L	<b>14.0</b>	<b>8.7</b>	<b>85.1</b>
(f5)	Hubert B	✓	Hubert B	18.8	11.1	82.3
(g1)	WavLM L			9.9	9.8	52.1
(g2)	WavLM L	✓	N/A	7.8	9.8	86.6
(g3)	WavLM L	✓	WavLM B+	7.2	<b>9.3</b>	88.3
(g4)	WavLM L	✓	WavLM L	<b>6.1</b>	10.1	<b>89.2</b>

\*Bold values indicate the best performing system.

understand tasks while maintaining the human perceptual quality, and thus it increases the generalizability of the SE for multiple tasks. Although the proposed system (c2) improved the WER of the Whisper ASR model by 21% compared with the baseline system (b), it still performs worse than the system without SE (a2).

System (c3) demonstrates the combination of SSL-MSE loss and OA with the addition ratio  $\beta = 0.1$ . The addition of the observed signal demonstrates the best SSL downstream task performance and the Whisper ASR performance among all systems. Improving the performance of a noise-robust ASR model with an SE model trained separately from ASR model is generally challenging (Kinoshita et al., 2020). Despite that, the proposed SE front-end in the system (c3) reduces the WER for the Whisper ASR model, which is trained to be robust to noise. In summary, the combination of OA with SSL-MSE loss training can further improve specific downstream task performance. Note that since the OA can be applied as a post-processing, the decision of whether or not to apply OA can change depending on the type of the downstream tasks. Thus, we can omit OA for human listening tasks using the same SE model.

While the OA technique alone, as used in system (b2), already shows promising generalization across downstream tasks without

requiring complex training procedures, a comparison with system (c2), which uses SSL-MSE loss without OA, reveals that the proposed SSL-MSE training provides consistent performance gains. Specifically, system (c2) achieves better results in human perceptual quality (DNSMOS OVRL: 3.42 vs. 2.93), SSL-based ASR (WER: 11.3% vs. 12.2%), and Whisper ASR (WER: 9.6% vs. 10.3%) compared to system (b2). These results indicate that, although OA is effective as a lightweight post-processing method, the SSL-MSE loss enables the SE model itself to produce higher-quality outputs that are more suitable for both human perception and machine-based downstream tasks. This suggests that the benefit of SSL-based training is not merely complementary to OA but essential for achieving a more universally robust enhancement system.

### 5.2.2 Comparison with other training criteria

Figure 3 shows the performance of SE systems trained with vanilla SNR loss, ASR loss, LMFB loss, and SSL-MSE loss. Each graph shows an evaluation metric: SE performance in SDR [dB], objective measures of perceptual quality in DNSMOS OVRL, SSL downstream ASR task performance in WER [%], and Whisper ASR performance in WER [%]. WavLM BASE+ is used to evaluate the SSL

TABLE 3 Effect of SE and SSL-MSE loss training on a noise-robust downstream model. We adopt WavLM Base plus model for SSL-MSE loss training.

	SE method	is downstream model noise robust	ASR WER↓ [%]	ASV EER↓ [%]	IC Acc↑ [%]
(a2)	no process		17.1	10.9	67.4
(h1)	no process	✓	13.6	9.2	79.4
(h2)	SNR loss	✓	12.5	7.9	85.8
(h3)	SSL-MSE loss	✓	<b>11.3</b>	<b>7.7</b>	<b>88.3</b>

\*Bold values indicate the best performing system.

downstream ASR performance. For each graph, the X-axis represents the multitask loss weight  $\alpha$  in the multitask training, which was fixed at  $\alpha = 0.1$  in the Table 1. The Y-axis indicates the performance. As for (a) SE performance and (b) objective measures of perceptual quality, all the multitask-trained systems share the trend that the performance approaches baseline SNR when multitask loss weight  $\alpha$  is large, while the performance drops when the multitask loss weight  $\alpha$  is small. The performance drop is more pronounced with ASR loss, compared with SSL-MSE loss and LMFB loss. This degradation is likely caused by the SE model becoming overly specialized for the ASR task.

(c) SSL-based ASR performance and (d) Whisper ASR performance are both improved by incorporating SSL-MSE, ASR loss, and LMFB loss, compared with SNR loss. Among these, the SSL-MSE loss yields the most substantial improvement, followed by the ASR loss. Although the ASR loss is task-matched, it does not provide the greatest performance gains. Compared to the Whisper ASR performance without SE (9.3%) shown in Table 1 (a2), the ASR loss actually degrades Whisper ASR performance across all multitask loss weight settings  $\alpha$ . This degradation appears to result from the SE model becoming overly specialized to the specific ASR model used during ASR-loss training, which differs from the Whisper ASR model, as described in Section 5.1.2. In other words, the ASR loss leads to insufficient generalization not only across different tasks, but also across different ASR models. Therefore, SSL-MSE loss improves SE generalizability to different tasks better than other common training criteria of ASR loss and LMFB loss.

The effect of the SSL-MSE loss on (c) SSL downstream ASR and (d) Whisper ASR performance is maximized when the multitask loss weight  $\alpha$  is set to 0.01. In contrast, (a) the SDR and (b) DNSMOS OVRL scores increase with larger values of  $\alpha$ , indicating a trade-off between human perceptual quality and downstream task performance.

### 5.2.3 Generalizability to other types of SSL models

Table 2 shows the SSL downstream task performances for a variation of a combination of SSL models that is used to calculate SSL-MSE loss and that is used for the inference of SSL downstream tasks.

Overall, WavLM BASE+ and WavLM LARGE demonstrate superior performance compared to other SSL models as the SSL models for inference. This is likely because these models are pre-trained with noise-aware strategies, making them more effective on test sets with additive noise. For the SSL pipeline with WavLM BASE+, SSL-MSE loss calculated with WavLM LARGE (c2) most

improves the performance, and SSL-MSE losses calculated with wav2vec 2.0 BASE (d1) and Hubert BASE (d2) do not improve the performance compared with vanilla SNR loss training (b). Moreover, SSL-MSE losses calculated with wav2vec 2.0 BASE and Hubert BASE do not improve matched condition where the same SSL model is used for inference, which is shown in (e5) and (f5). On the other hand, SSL-MSE loss calculated with WavLM BASE+ and WavLM LARGE improves every SSL pipeline with different types of SSL models including WavLM BASE+ shown in (c1) and (c2), Wav2vec 2.0 BASE shown in (e3) and (e4), Hubert BASE shown in (f3) and (f4), and WavLM LARGE (g3) and (g4), compared with the SE model trained with SNR loss. Thus, it can be said that SSL-MSE training potentially generalizes over different types of SSL models. Compared with SSL-MSE loss calculated with WavLM BASE+, WavLM LARGE offers better performance for the SSL pipeline with different SSL models. The higher generalizability of WavLM LARGE compared with WavLM BASE+ as the upstream model (Chen et al., 2022) seems to transfer more general knowledge to the SE model via SSL-MSE loss training.

### 5.2.4 The effect of SSL-MSE loss on noise-robust downstream model

To further discuss the effectiveness of SE and SSL-MSE, we prepare an SSL pipeline where the downstream model is also trained to be robust to noise by the noisy paired data for each task. Table 3 shows the SSL pipeline performance for three downstream tasks. Since WavLM BASE+ is used for the SSL pipeline, SSL upstream itself is also robust to noise. The table shows that the noise-robust training (h1) improved the performance of each task compared to the system without noise-robust training (a2). This is notable that the introduction of baseline SE front-end (h2) further improves these performances despite the fact that both SSL upstream and downstream are trained to be robust to noise. The introduction of the proposed SSL-MSE loss (h3) substantially improves the performance of system (h2). This indicates that SSL-MSE loss is effective for the back-end tasks that are trained to be robust to noise.

## 6 Conclusion

To build a general speech enhancement (SE) model that can be applied to various downstream tasks, we proposed a self-supervised learning mean squared error (SSL-MSE) loss training scheme that exploits the rich speech representation captured by the SSL models

to train SE models. By calculating the loss term on the feature domain of the SSL model, SE is expected to be trained to preserve high-level information contained in speech audio that is crucial to solving various downstream tasks. The experimental validation shows that introducing SSL-MSE loss improves SSL downstream task performance and Whisper ASR performance while maintaining DNSMOS results. In addition to this, the combination of OA with SSL-MSE further improves the performance of back-end tasks. Further analyses show that SSL-MSE loss can generalize over other types of SSL models than that used to calculate SSL-MSE loss. It is also shown that SSL-MSE loss training requires the use of a noise-robust SSL model like WavLM BASE+ to improve performance. Moreover, SSL-MSE is effective in cases where both the upstream SSL model and downstream task model are trained to be robust to noise.

This study identifies several directions for future research. First, it is important to validate the effectiveness of the proposed approach across a wider range of downstream tasks. In this work, we demonstrated the generalizability of the model trained with SSL-MSE by evaluating its performance on three types of downstream tasks: human listening tests, tasks using SSL models, and Whisper ASR. As a future direction, we plan to explore its applicability to additional tasks such as off-the-shelf voice activity detection (VAD) (Wiseman, 2025), speaker diarization models (Bredin, 2023), etc. Second, the evaluation of objective measures of perceptual quality using actual subjective listening tests remains an essential future task. While this study employed DNSMOS to assess the perceptual quality and demonstrated that the proposed method can improve both perceptual quality and machine understanding tasks, conducting subjective listening tests with human listeners will enable a more reliable assessment. Third, it is necessary to investigate the applicability of the proposed approach to a broader set of speech enhancement models with different characteristics. In this work, we performed experimental validation using Conv-TasNet, a widely adopted model in the field. Future work will include exploring the generalizability of the proposed approach to various real-time speech enhancement methods (Hu et al., 2020) and larger-scale speech enhancement models (Wang et al., 2023).

## Data availability statement

The datasets presented in this article are not readily available because The dataset in this research is generated based on open

dataset and the way to generate the data is described in the manuscript. However, the generated dataset itself cannot be shared due to corporate restrictions. Requests to access the datasets should be directed to hrs.sato@ntt.com.

## Author contributions

HS: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Project administration, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review and editing. TO: Methodology, Supervision, Writing – review and editing. MD: Methodology, Supervision, Writing – review and editing. TM: Supervision, Writing – review and editing. TA: Writing – review and editing. RM: Conceptualization, Writing – review and editing.

## Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

## Conflict of interest

Authors HS, TO, MD, TM, TA, and RM were employed by NTT Corporation.

## Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## References

- Ali, M. N., Brutti, A., and Falavigna, D. (2023). Direct enhancement of pre-trained speech embeddings for speech processing in noisy conditions. *Comput. Speech and Lang.* 81, 101501. doi:10.1016/j.csl.2023.101501
- Ali, M. N., Brutti, A., and Falavigna, D. (2022). Enhancing embeddings for speech classification in noisy conditions. *Proc. Interspeech*, 2933–2937. doi:10.21437/Interspeech.2022-10707
- Avila, A. R., Alam, M. J., O'Shaughnessy, D., and Falk, T. (2018). Investigating speech enhancement and perceptual quality for speech emotion recognition. *Proc. Interspeech*, 3663–3667. doi:10.21437/Interspeech.2018-2350
- Baevski, A., Schneider, S., and Auli, M. (2020a). “vq-wav2vec: self-supervised learning of discrete speech representations,” in *8th international conference on learning representations (OpenReview.net)*.
- Baevski, A., Zhou, Y., Mohamed, A., and Auli, M. (2020b). wav2vec 2.0: a framework for self-supervised learning of speech representations. *Proc. Adv. neural Inf. Process. Syst.* 33, 12449–12460.
- Bredin, H. (2023). pyannote. audio 2.1 speaker diarization pipeline: principle, benchmark, and recipe. *Proc. Interspeech*, 1983–1987. doi:10.21437/Interspeech.2023-105
- Chang, H.-J., Yang, S.-W., and Lee, H.-Y. (2022a). “Distillhubert: speech representation learning by layer-wise distillation of hidden-unit bert,” in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 7087–7091.
- Chang, X., Maekaku, T., Fujita, Y., and Watanabe, S. (2022b). “End-to-End integration of speech recognition, speech enhancement, and self-supervised learning representation,” in *Proc. Annual conference of the international speech communication association*, 3819–3823.
- Chen, S., Wang, C., Chen, Z., Wu, Y., Liu, S., Chen, Z., et al. (2022). Wavlm: large-scale self-supervised pre-training for full stack speech processing. *IEEE J. Sel. Top. Signal Process.* 16, 1505–1518. doi:10.1109/jstsp.2022.3188113

- Chen, S.-J., Subramanian, A. S., Xu, H., and Watanabe, S. (2018). "Building state-of-the-art distant speech recognition using the chime-4 challenge with a setup of speech enhancement baseline," in *Proc. Annual conference of the international speech communication association*, 1571–1575.
- Chen, W., Shi, J., Yan, B., Berrebbi, D., Zhang, W., Peng, Y., et al. (2023). "Joint prediction and denoising for large-scale multilingual self-supervised learning," in *IEEE automatic speech Recognition and understanding workshop (IEEE)*, 1–8.
- Défossez, A., Synnaeve, G., and Adi, Y. (2020). "Real time speech enhancement in the waveform domain," in *Proc. Annual conference of the international speech communication association*, 3291–3295.
- Delcroix, M., Ochiai, T., Zmolikova, K., Kinoshita, K., Tawara, N., Nakatani, T., et al. (2020). "Improving speaker discrimination of target speech extraction with time-domain SpeakerBeam," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 691–695.
- Delcroix, M., Žmoliková, K., Kinoshita, K., Ogawa, A., and Nakatani, T. (2018). "Single channel target speaker extraction and recognition with speaker beam," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 5554–5558.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). "BERT: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 conference of the north American chapter of the association for computational linguistics: human language technologies* (Association for Computational Linguistics), 1, 4171–4186.
- Dubey, H., Aazami, A., Gopal, V., Naderi, B., Braun, S., Cutler, R., et al. (2024). Icassp 2023 deep noise suppression challenge. *IEEE Open J. Signal Process.* 5, 725–737. doi:10.1109/ojsp.2024.3378602
- Dubey, H., Gopal, V., Cutler, R., Aazami, A., Matushevych, S., Braun, S., et al. (2022). "ICASSP 2022 deep noise suppression challenge," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 9271–9275.
- Dunbar, E., Bernard, M., Hamilakis, N., Nguyen, T. A., de Seyssel, M., Rozé, P., et al. (2021). "The zero resource speech challenge 2021: spoken Language modelling," in *Proc. Annual conference of the international speech communication association*, 1574–1578.
- Eskimez, S. E., Soufleris, P., Duan, Z., and Heinzelman, W. (2018). Front-end speech enhancement for commercial speaker verification systems. *Speech Commun.* 99, 101–113. doi:10.1016/j.specom.2018.03.008
- Fujimoto, M., and Kawai, H. (2019). "One-pass single-channel noisy speech recognition using a combination of noisy and enhanced features," in *Proc. Annual conference of the international speech communication association*, 486–490.
- Fujita, K., Sato, H., Ashihara, T., Kanagawa, H., Delcroix, M., Moriya, T., et al. (2024). "Noise-robust zero-shot text-to-speech synthesis conditioned on self-supervised speech-representation model with adapters," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 11471–11475.
- Gao, T., Du, J., Dai, L.-R., and Lee, C.-H. (2015). "Joint training of front-end and back-end deep neural networks for robust speech recognition," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 4375–4379.
- GitHub (2024). s3prl toolkit [Dataset]. Available online at: <https://github.com/s3prl/s3prl> (Accessed October 7, 2024).
- Graves, A., Fernandez, S., Gomez, F., and Schmidhuber, J. (2006). "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *International conference on machine learning (PMLR)*, 369–376.
- Hao, X., Su, X., Horaud, R., and Li, X. (2021). "Fullsubnet: a full-band and sub-band fusion model for real-time single-channel speech enhancement," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 6633–6637.
- Hsieh, T.-A., Yu, C., Fu, S.-W., Lu, X., and Tsao, Y. (2021). "Improving perceptual quality by phone-fortified perceptual loss using wasserstein distance for speech enhancement," in *Proc. Annual conference of the international speech communication association*, 196–200.
- Hsu, W.-N., Bolte, B., Tsai, Y.-H. H., Lakhotia, K., Salakhutdinov, R., and Mohamed, A. (2021). Hubert: self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Trans. Audio, Speech, Lang. Process.* 29, 3451–3460. doi:10.1109/taslp.2021.3122291
- Hu, Y., Liu, Y., Lv, S., Xing, M., Zhang, S., Fu, Y., et al. (2020). "DCCRN: deep complex convolution recurrent network for phase-aware speech enhancement," in *Proc. Annual conference of the international speech communication association*, 2472–2476.
- Huh, J., Brown, A., Jung, J.-w., Chung, J. S., Nagrani, A., Garcia-Romero, D., et al. (2023). Voxsrc 2022: the fourth voxceleb speaker recognition challenge. *arXiv preprint arXiv:2302.10248*
- Iwamoto, K., Ochiai, T., Delcroix, M., Ikeshita, R., Sato, H., Araki, S., et al. (2022). "How bad are artifacts? Analyzing the impact of speech enhancement errors on ASR," in *Proc. Annual conference of the international speech communication association*, 5418–5422.
- Jaiswal, A., Babu, A. R., Zadeh, M. Z., Banerjee, D., and Makedon, F. (2020). A survey on contrastive self-supervised learning. *Technologies* 9 (2), 2. doi:10.3390/technologies9010002
- Ju, Y., Zhang, S., Rao, W., Wang, Y., Yu, T., Xie, L., et al. (2023). "Tea-pse 2.0: sub-band network for real-time personalized speech enhancement," in *Proc. SLT 2022 (IEEE)*, 472–479.
- Kandagatla, R. K., and Potluri, V. S. (2020). Performance analysis of neural network, nmf and statistical approaches for speech enhancement. *Int. J. Speech Technol.* 23, 917–937. doi:10.1007/s10772-020-09751-6
- Karita, S., Soplin, N. E. Y., Watanabe, S., Delcroix, M., Ogawa, A., and Nakatani, T. (2019). "Improving transformer-based end-to-end speech recognition with connectionist temporal classification and language model integration," in *Proc. Annual conference of the international speech communication association*, 1408–1412.
- Kingma, D. P., and Ba, J. (2015). "Adam: a method for stochastic optimization," in *Proc. International conference on learning representations*.
- Kinoshita, K., Ochiai, T., Delcroix, M., and Nakatani, T. (2020). "Improving noise robust automatic speech recognition with single-channel time-domain enhancement network," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 7009–7013.
- Le Roux, J., Wisdom, S., Erdogan, H., and Hershey, J. R. (2019). "Sdr-half-baked or well done?," in *Proc. IEEE international Conference on acoustics, speech and signal processing (IEEE)*, 626–630.
- Liu, A. T., Li, S.-W., and yi Lee, H. (2020a). Tera: self-supervised learning of transformer encoder representation for speech [Dataset]. doi:10.48550/arXiv.2007.06028
- Liu, A. T., Yang, S.-W., Chi, P.-H., Hsu, P.-C., and Lee, H.-Y. (2020b). "Mockingjay: unsupervised speech representation learning with deep bidirectional transformer encoders," in *Proc. IEEE international conference on acoustics, speech and signal processing*. doi:10.1109/icassp40776.2020.9054458
- Liu, X., Zhang, F., Hou, Z., Mian, L., Wang, Z., Zhang, J., et al. (2021). Self-supervised learning: generative or contrastive. *IEEE Trans. Knowl. Data Eng.* 35, 857–876. doi:10.1109/TKDE.2021.3090866
- Lu, Y.-J., Chang, X., Li, C., Zhang, W., Cornell, S., Ni, Z., et al. (2022). "ESPnet-SE++: speech enhancement for robust speech recognition, translation, and understanding," in *Proc. Annual conference of the international speech communication association*, 5458–5462.
- Luo, Y., and Mesgarani, N. (2019). Conv-TasNet: surpassing ideal time-frequency magnitude masking for speech separation. *IEEE/ACM Trans. Audio, Speech, Lang. Process.* 27, 1256–1266. doi:10.1109/taslp.2019.2915167
- Masuyama, Y., Chang, X., Cornell, S., Watanabe, S., and Ono, N. (2023). "End-to-end integration of speech recognition, dereverberation, beamforming, and self-supervised learning representation," in *Proc. IEEE spoken language technology workshop (IEEE)*, 260–265.
- Menne, T., Schlüter, R., and Ney, H. (2019). "Investigation into joint optimization of single channel speech enhancement and acoustic modeling for robust asr," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 6660–6664.
- Mohamed, A., Lee, H.-y., Borgholt, L., Havtorn, J. D., Edin, J., Igel, C., et al. (2022). Self-supervised speech representation learning: a review. *IEEE J. Sel. Top. Signal Process.* 16, 1179–1210. doi:10.1109/jstsp.2022.3207050
- Ochiai, T., Iwamoto, K., Delcroix, M., Ikeshita, R., Sato, H., Araki, S., et al. (2024). Rethinking processing distortions: disentangling the impact of speech enhancement errors on speech recognition performance. *IEEE/ACM Trans. Audio Speech Lang. Process.* 32, 3589–3602. doi:10.1109/taslp.2024.3426924
- O'Malley, T., Narayanan, A., Wang, Q., Park, A., Walker, J., and Howard, N. (2021). "A conformer-based asr frontend for joint acoustic echo cancellation, speech enhancement and speech separation," in *Proc. IEEE automatic speech Recognition and understanding workshop (IEEE)*, 304–311.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. (2015). "LibriSpeech: an ASR corpus based on public domain audio books," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 5206–5210.
- Pandey, A., Liu, C., Wang, Y., and Saraf, Y. (2021). "Dual application of speech enhancement for automatic speech recognition," in *Proc. SLT 2021 (IEEE)*, 223–228.
- Pandey, A., and Wang, D. (2020). "Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 6629–6633.
- Pasad, A., Chou, J.-C., and Livescu, K. (2021). "Layer-wise analysis of a self-supervised speech representation model," in *Proc. IEEE automatic speech Recognition and understanding workshop (IEEE)*, 914–921.
- Prabhavalkar, R., Hori, T., Sainath, T. N., Schlüter, R., and Watanabe, S. (2023). "End-to-end speech recognition: a survey," in *IEEE/ACM transactions on audio, speech, and language processing*.
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. (2023). "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning (PMLR)*, 28492–28518.
- Reddy, C. K., Dubey, H., Koishida, K., Nair, A., Gopal, V., Cutler, R., et al. (2021). "INTERSPEECH 2021 deep noise suppression challenge," in *Proc. Annual conference of the international speech communication association*, 2796–2800.
- Reddy, C. K., Gopal, V., and Cutler, R. (2022). "Dnsmos p. 835: a non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 886–890.



- Reddy, C. K., Gopal, V., Cutler, R., Beyrami, E., Cheng, R., Dubey, H., et al. (2020). The interspeech 2020 deep noise suppression challenge: datasets, subjective testing framework, and challenge results. *arXiv Prepr. arXiv:2005.13981*.
- Rix, A. W., Beerends, J. G., Hollier, M. P., and Hekstra, A. P. (2001). "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)* (IEEE), 2, 749–752. doi:10.1109/icassp.2001.941023
- Sato, H., Masumura, R., Ochiai, T., Delcroix, M., Moriya, T., Ashihara, T., et al. (2023). "Downstream task agnostic speech enhancement with self-supervised representation loss," in *Proc. Annual conference of the international speech communication association*, 854–858.
- Sato, H., Moriya, T., Mimura, M., Horiguchi, S., Ochiai, T., Ashihara, T., et al. (2024). Speakerbeam-ss: real-time target speaker extraction with lightweight conv-tasnet and state space modeling. *arXiv Prepr. arXiv:2407.01857*, 5033–5037. doi:10.21437/interspeech.2024-413
- Sato, H., Ochiai, T., Delcroix, M., Kinoshita, K., Kamo, N., and Moriya, T. (2022). "Learning to enhance or not: neural network-based switching of enhanced and observed signals for overlapping speech recognition," in *Proc. IEEE international conference on acoustics, speech and signal processing*, 6287–6291.
- Sato, H., Ochiai, T., Delcroix, M., Kinoshita, K., Moriya, T., and Kamo, N. (2021). "Should we always separate? Switching between enhanced and observed signals for overlapping speech recognition," in *Proc. Annual conference of the international speech communication association*, 1149–1153.
- Schneider, S., Baevski, A., Collobert, R., and Auli, M. (2019). "wav2vec: unsupervised pre-training for speech recognition," in *Proc. Annual conference of the international speech communication association*, 3465–3469.
- Shi, J., Inaguma, H., Ma, X., Kulikov, I., and Sun, A. Y. (2024). "Multi-resolution hubert: multi-resolution speech self-supervised learning with masked unit prediction," in *The twelfth international Conference on learning representations, ICLR 2024, Vienna, Austria, may 7–11, 2024 (OpenReview.net)*.
- Shon, S., Tang, H., and Glass, J. (2019). "VoiceID loss: speech enhancement for speaker verification," in *Proc. Annual conference of the international speech communication association*, 2888–2892.
- Speech to text (2024). Speech to text - OpenAI API [dataset]. Available online at: <https://platform.openai.com/docs/guides/speech-to-text> (Accessed October 7, 2024).
- Speech-to-Text AI (2024). Speech-to-Text AI [dataset]. Available online at: <https://cloud.google.com/speech-to-text> (Accessed October 7, 2024).
- Subramanian, A. S., Wang, X., Baskar, M. K., Watanabe, S., Taniguchi, T., Tran, D., et al. (2019). "Speech enhancement using end-to-end speech recognition objectives," in *Proc. WASPAA 2019 (IEEE)*, 234–238.
- Thakker, M., Eskimez, S. E., Yoshioka, T., and Wang, H. (2022). "Fast real-time personalized speech enhancement: end-to-end enhancement network (E3Net) and knowledge distillation," in *Proc. Annual conference of the international speech communication association*, 991–995.
- Tsai, H.-S. (2022). "SUPERB-SG: enhanced speech processing universal PERFORMANCE benchmark for semantic and generative capabilities," in *Proceedings of the 60th annual meeting of the association for computational linguistics (ACL)*, 8479–8492.
- Vincent, E., Gribonval, R., and Févotte, C. (2006). Performance measurement in blind audio source separation. *IEEE Trans. Audio, Speech, Lang. Process.* 14, 1462–1469. doi:10.1109/tsa.2005.858005
- Wang, K., He, B., and Zhu, W.-P. (2021). "Tstnn: two-stage transformer based neural network for speech enhancement in the time domain," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 7098–7102.
- Wang, Q., Moreno, I. L., Saglam, M., Wilson, K., Chiao, A., Liu, R., et al. (2020). "VoiceFilter-Lite: streaming targeted voice separation for on-device speech recognition," in *Proc. Annual conference of the international speech communication association*, 2677–2681.
- Wang, Q., Muckenhirn, H., Wilson, K., Sridhar, P., Wu, Z., Hershey, J. R., et al. (2019). "VoiceFilter: targeted voice separation by speaker-conditioned spectrogram masking," in *Proc. Annual conference of the international speech communication association*, 2728–2732.
- Wang, Z.-Q., Cornell, S., Choi, S., Lee, Y., Kim, B.-Y., and Watanabe, S. (2023). "Tf-gridnet: making time-frequency domain models great again for monaural speaker separation," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 1–5.
- Wang, Z.-Q., and Wang, D. (2016). A joint training framework for robust automatic speech recognition. *IEEE/ACM Trans. Audio, Speech, Lang. Process.* 24, 796–806. doi:10.1109/taslp.2016.2528171
- Wani, T. M., Gunawan, T. S., Qadri, S. A. A., Kartiwi, M., and Ambikairajah, E. (2021). A comprehensive review of speech emotion recognition systems. *IEEE access* 9, 47795–47814. doi:10.1109/access.2021.3068045
- Watanabe, S., Hori, T., Karita, S., Hayashi, T., Nishitoba, J., Unno, Y., et al. (2018). "Espnet: end-to-end speech processing toolkit," in *Proc. Annual conference of the international speech communication association*, 2207–2211.
- Watanabe, S., Hori, T., Kim, S., Hershey, J., and Hayashi, T. (2017). Hybrid CTC/Attention architecture for end-to-end speech recognition. *IEEE J. Sel. Top. Signal Process.* 11, 1240–1253. doi:10.1109/jstsp.2017.2763455
- Wiseman (2025). WebRTC voice activity detector [dataset]. Available online at: <https://github.com/wiseman/py-webrtcvad> (Accessed June 1, 2025).
- Xia, Y., Braun, S., Reddy, C. K., Dubey, H., Cutler, R., and Tashev, I. (2020). "Weighted speech distortion losses for neural-network-based real-time speech enhancement," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 871–875.
- Yadav, H., Sitaram, S., and Shah, R. R. (2024). "Ms-hubert: mitigating pre-training and inference mismatch in masked language modelling methods for learning speech representations," in *Proc. Annual conference of the international speech communication association*, 5053–5057.
- Yang, S.-w., Chi, P.-H., Chuang, Y.-S., Lai, C.-I. J., Lakhotia, K., Lin, Y. Y., et al. (2021). "SUPERB: speech processing universal PERFORMANCE benchmark," in *Proc. Annual conference of the international speech communication association*, 1194–1198.
- Yoshioka, T., Ito, N., Delcroix, M., Ogawa, A., Kinoshita, K., Fujimoto, M., et al. (2015). "The ntt chime-3 system: advances in speech enhancement and recognition for mobile multi-microphone devices," in *Proc. IEEE workshop on automatic speech recognition and understanding*, 436–443.
- Zhao, S., Ma, B., Watcharasupat, K. N., and Gan, W.-S. (2022). "Frcrn: boosting feature representation using frequency recurrence for monaural speech enhancement," in *Proc. IEEE international conference on acoustics, speech and signal processing (IEEE)*, 9281–9285.
- Žmolíková, K., Delcroix, M., Kinoshita, K., Ochiai, T., Nakatani, T., Burget, L., et al. (2019). SpeakerBeam: speaker aware neural network for target speaker extraction in speech mixtures. *IEEE J. Sel. Top. Signal Process.* 13, 800–814. doi:10.1109/jstsp.2019.2922820